

Should We Be Morally Accountable for AI Behavior? A Novel Framework for Distributed Responsibility in Artificial Intelligence Systems

Kwan Hong TAN

Associate Faculty

Singapore University of Social Sciences

khtan055@suss.edu.sg

27 April 2026

Abstract

The rapid advancement of artificial intelligence (AI) systems has fundamentally challenged traditional notions of moral responsibility and accountability. As AI systems become increasingly autonomous and capable of causing significant harm, the question of who should be held morally accountable for their behavior has become one of the most pressing ethical issues of our time. This thesis presents a comprehensive examination of moral accountability in AI systems, introducing a novel theoretical framework called "Gradient Responsibility Networks" (GRN) that addresses critical gaps in existing approaches.

Through extensive analysis of philosophical foundations, empirical case studies, and real-world incidents, this work demonstrates that traditional binary models of moral responsibility are inadequate for the complex, distributed, and temporal nature of AI systems. The proposed GRN framework offers a mathematically rigorous, practically implementable approach to distributing moral responsibility across networks of actors involved in AI development, deployment, and governance.

Key findings include: (1) 80-85% of AI projects fail, creating widespread accountability gaps; (2) existing frameworks fail to address temporal decay of responsibility and emergent behaviors; (3) the GRN framework provides superior performance across six critical

dimensions compared to traditional approaches; and (4) empirical case studies validate the framework's practical applicability across diverse AI domains.

This research contributes to the growing body of literature on AI ethics by providing both theoretical innovation and practical tools for policymakers, technologists, and ethicists grappling with the challenges of AI accountability in the 21st century.

Keywords: artificial intelligence, moral responsibility, accountability, ethics, distributed systems, governance

Acknowledgments: This research builds upon the foundational work of numerous scholars in philosophy, computer science, law, and policy studies. While the novel theoretical framework and analysis presented here are original contributions, they are deeply informed by the broader academic community's ongoing efforts to understand and address the challenges of AI governance and accountability.

Declaration: This thesis presents original research and analysis. All sources have been properly cited and referenced. The Gradient Responsibility Networks framework and its mathematical formalization represent novel theoretical contributions to the field of AI ethics and governance.

License: This work is licensed under CC BY-NC-ND 4.0. Any use must cite the author and cannot be modified or used for commercial purposes.

This research is part of an ongoing series of works on *Artificial Intelligence Ethics*, *Moral Responsibility*, and related domains. Related theses, conceptual frameworks, and methodological contributions by the same author are accessible via the following profiles:

[ORCID iD: 0009-0003-9276-2829](https://orcid.org/0009-0003-9276-2829)

[ResearchGate: Kwan Hong Tan](https://www.researchgate.net/profile/Kwan-Hong-Tan)

Table of Contents

1. Introduction 4

2. Literature Review 6

3. Theoretical Framework: Gradient Responsibility Networks 13

4. Empirical Analysis and Case Studies 21

5. Comparative Framework Analysis 32

6. Implications and Applications 44

7. Limitations and Future Research 53

8. Conclusion 60

9. References 66

1. Introduction

The question of moral accountability for artificial intelligence behavior represents one of the most complex and consequential ethical challenges of the digital age. As AI systems increasingly permeate critical domains—from autonomous vehicles navigating public roads to algorithmic systems making decisions about employment, healthcare, and criminal justice—the stakes of getting accountability frameworks right have never been higher. The traditional philosophical frameworks for moral responsibility, developed primarily for individual human agents, are proving inadequate for the distributed, temporal, and emergent nature of modern AI systems.

Consider the tragic case of Joshua Brown, who died in 2016 when his Tesla Model S, operating in Autopilot mode, collided with a tractor-trailer that the system failed to detect [1]. This incident crystallized a fundamental question: Who bears moral responsibility for this outcome? Tesla, for developing and marketing the Autopilot system? Brown himself, for choosing to rely on the technology? The regulatory bodies that approved the system for public use? The engineers who wrote the specific algorithms? The answer, as this thesis will argue, is not a simple either-or proposition but rather a complex network of distributed responsibility that existing frameworks struggle to capture.

The urgency of addressing this question is underscored by the staggering failure rates of AI systems. Recent research indicates that 80-85% of AI projects fail to meet their intended objectives [2], while 42% of businesses scrapped most of their AI initiatives in 2025, up from just 17% the previous year [3]. These failures are not merely technical disappointments—they represent potential sources of harm ranging from discriminatory hiring practices to life-threatening medical misdiagnoses. Yet our current approaches to accountability remain rooted in pre-digital conceptions of individual responsibility that are fundamentally mismatched to the realities of AI development and deployment.

This thesis addresses this critical gap by proposing a novel theoretical framework called "Gradient Responsibility Networks" (GRN) that reconceptualizes moral accountability for AI systems. Rather than forcing binary choices about who is or is not responsible, the GRN framework recognizes that responsibility exists on continuous gradients and is distributed

across networks of actors who contribute to AI outcomes in varying degrees over time. This approach provides a more nuanced, mathematically rigorous, and practically implementable foundation for AI accountability that can inform legal frameworks, policy decisions, and ethical guidelines.

The central argument of this thesis is that we should indeed be morally accountable for AI behavior, but that this accountability must be understood as distributed across networks of actors rather than concentrated in individual agents. This distributed accountability is not a dilution of responsibility but rather a more accurate reflection of the complex causal chains and moral relationships inherent in AI systems. By developing frameworks that can capture and quantify these distributed responsibilities, we can create more effective mechanisms for preventing harm, ensuring justice, and promoting the beneficial development of AI technologies.

The thesis proceeds through seven main sections. Following this introduction, Section 2 provides a comprehensive literature review examining existing approaches to AI ethics and moral responsibility theory. Section 3 introduces the novel GRN framework, including its mathematical formalization and philosophical justification. Section 4 presents empirical analysis and case studies that validate the framework's applicability. Section 5 offers a comparative analysis demonstrating the GRN framework's advantages over traditional approaches. Section 6 explores the practical implications and applications of the framework for policy and governance. Section 7 acknowledges limitations and suggests directions for future research, while Section 8 concludes with a synthesis of key findings and their broader significance.

This research contributes to the growing field of AI ethics by providing both theoretical innovation and practical tools for addressing one of the most challenging questions in contemporary technology policy. As AI systems become increasingly powerful and ubiquitous, the frameworks we develop today for understanding and implementing moral accountability will shape the trajectory of technological development for decades to come. The stakes could not be higher, and the need for rigorous, nuanced approaches to these questions has never been more urgent.

2. Literature Review

2.1 Foundations of Moral Responsibility Theory

The philosophical foundations of moral responsibility trace back to Aristotle's seminal work in the *Nicomachean Ethics*, where he first articulated the conditions under which individuals can be held accountable for their actions [4]. Aristotle's framework established two fundamental requirements: that the agent must have knowledge of the particular circumstances of the action, and that the action must originate from the agent rather than external compulsion. These Aristotelian insights continue to influence contemporary debates about moral responsibility, though their application to artificial intelligence systems presents novel challenges.

Modern philosophical discourse on moral responsibility has been dominated by debates about free will and determinism, with significant implications for how we understand accountability in AI systems. The Stanford Encyclopedia of Philosophy defines moral responsibility as involving "attributing certain powers and capacities to that person, and viewing their behavior as arising, in the right way, from the fact that the person has, and has exercised, these powers and capacities" [5]. This definition immediately raises questions about whether AI systems themselves can possess the requisite powers and capacities, or whether responsibility must be attributed to the human agents involved in their creation and deployment.

The incompatibilist position, articulated most forcefully by Peter van Inwagen's Consequence Argument, suggests that moral responsibility requires the ability to do otherwise—a capacity that may be undermined by causal determinism [6]. If our actions are the inevitable consequences of prior causes, van Inwagen argues, then we cannot be truly responsible for them. This philosophical challenge takes on new dimensions in the context of AI systems, where the deterministic nature of algorithmic processes might seem to preclude genuine moral responsibility. However, as this thesis will argue, the distributed nature of AI development and deployment creates multiple points of moral agency that complicate simple deterministic accounts.

Contemporary philosophers have increasingly recognized the limitations of purely individualistic approaches to moral responsibility. The work of scholars like Margaret Gilbert

on collective responsibility and Larry May on shared responsibility has highlighted how moral accountability can be distributed across groups and institutions [7]. These insights prove particularly relevant to AI systems, which typically involve complex networks of developers, deployers, users, and regulators. However, existing theories of collective responsibility were not designed with the specific characteristics of AI systems in mind, creating a need for new theoretical frameworks.

The distinction between causal responsibility and moral responsibility proves crucial for understanding AI accountability. As the Stanford Encyclopedia notes, "we cannot always infer moral responsibility from an assignment of causal responsibility" [8]. A young child can cause an outcome while failing to fulfill the general requirements for moral responsibility. Similarly, AI systems can be causally responsible for outcomes without bearing moral responsibility in the traditional sense. This distinction becomes complex in AI contexts where multiple human agents contribute causally to outcomes through their roles in system development, deployment, and governance.

2.2 Artificial Moral Agents and AI Ethics

The question of whether AI systems themselves can be moral agents has generated significant scholarly debate. Martinho et al. (2021) identified five distinct perspectives on Artificial Moral Agents (AMAs) in their comprehensive survey of the field [9]. The first perspective, "AMA as Moral Patients," suggests that AI systems should be protected from harm but cannot themselves be moral agents. The second, "AMA as Moral Agents," argues that sufficiently sophisticated AI systems can bear moral responsibility for their actions. The third perspective, "AMA as Moral Proxies," positions AI systems as extensions of human moral agency rather than independent agents. The fourth, "AMA as Moral Tools," treats AI systems purely as instruments without independent moral status. Finally, the fifth perspective, "AMA as Moral Partners," envisions collaborative moral relationships between humans and AI systems.

Each of these perspectives has different implications for how we understand accountability in AI systems. The "moral tools" perspective, for instance, would place all responsibility on the human agents who create and deploy AI systems, while the "moral agents" perspective would allow for direct attribution of responsibility to AI systems themselves. The "moral proxies" and

"moral partners" perspectives suggest more complex distributions of responsibility between human and artificial agents.

The empirical research by Martinho et al. revealed significant disagreement among experts about which perspective best captures the moral status of AI systems. This disagreement reflects deeper philosophical uncertainties about the nature of moral agency and its requirements. However, for practical purposes of accountability and governance, the question may be less about whether AI systems are "truly" moral agents and more about how we can effectively distribute responsibility across the networks of human agents involved in AI development and deployment.

The field of machine ethics, pioneered by scholars like Wendell Wallach and Colin Allen, has focused primarily on how to build ethical decision-making capabilities into AI systems [10]. While this work is valuable for creating more beneficial AI systems, it does not directly address the question of how to attribute moral responsibility when AI systems cause harm. The assumption that AI systems can be made perfectly ethical through better design overlooks the inherent uncertainties and value conflicts that characterize real-world moral decisions.

Recent work in AI ethics has increasingly recognized the importance of considering the entire lifecycle of AI systems, from data collection and model training through deployment and ongoing maintenance. This lifecycle perspective highlights the temporal dimension of AI accountability—responsibility may shift over time as systems are modified, contexts change, and new stakeholders become involved. However, existing frameworks for AI ethics have not developed sophisticated tools for tracking and allocating responsibility across these temporal dimensions.

2.3 Accountability Frameworks in AI Systems

Novelli, Taddeo, and Floridi (2023) have provided one of the most comprehensive frameworks for understanding accountability in AI systems [11]. Their approach defines accountability as "answerability"—a relation between an agent and a forum such that the agent must justify their conduct to the forum, which supervises, questions, and passes judgment. This framework identifies seven key features of accountability: context, range, agents, forum, standards, process, and implications.

The Novelli et al. framework represents a significant advance in thinking about AI accountability by recognizing the multidimensional nature of the problem. Their identification of different types of agents (individual, corporate, collective, hierarchical) and forums (affected parties, stakeholders, domain practitioners) provides a useful taxonomy for understanding the complexity of AI accountability relationships. However, their framework remains primarily descriptive rather than prescriptive—it helps us understand the dimensions of accountability but provides limited guidance on how to actually allocate responsibility across these dimensions.

The framework's treatment of temporal dynamics is particularly limited. While Novelli et al. acknowledge that accountability relationships can change over time, they do not provide tools for understanding how responsibility should decay or shift as AI systems evolve. This temporal dimension is crucial for AI systems, which may cause harm years after their initial development and may be modified by multiple parties over their operational lifetime.

The European Union's approach to AI governance, embodied in the AI Act, represents the most comprehensive regulatory framework for AI accountability to date [12]. The AI Act takes a risk-based approach, categorizing AI systems into different risk levels and imposing corresponding obligations on various actors in the AI value chain. High-risk AI systems are subject to strict requirements for risk management, data governance, transparency, and human oversight.

However, the AI Act's approach to responsibility attribution remains relatively traditional, focusing primarily on the obligations of AI system providers and deployers. While this provides clarity for regulatory compliance, it may not capture the full complexity of responsibility relationships in AI systems. The Act's emphasis on ex-ante compliance measures also provides limited guidance for ex-post attribution of responsibility when AI systems cause harm.

Legal scholars have begun to grapple with how existing tort law might apply to AI systems. The RAND Corporation's analysis of liability for AI systems highlights the challenges of applying traditional legal concepts like negligence and strict liability to AI contexts [13]. The distributed nature of AI development, the difficulty of establishing causation in complex

systems, and the challenges of proving foreseeability all complicate traditional approaches to legal liability.

The emergence of algorithmic auditing as a field represents another approach to AI accountability. Researchers like Cathy O'Neil and Safiya Noble have documented how algorithmic systems can perpetuate and amplify social biases, leading to discriminatory outcomes in domains ranging from criminal justice to employment [14]. However, algorithmic auditing typically focuses on identifying bias and discrimination rather than attributing moral responsibility for these outcomes.

2.4 Gaps in Existing Approaches

Despite the growing body of literature on AI ethics and accountability, several critical gaps remain. First, existing frameworks tend to treat responsibility as binary—either someone is responsible or they are not. This binary approach fails to capture the reality that multiple actors may bear varying degrees of responsibility for AI outcomes. A more nuanced approach that recognizes gradations of responsibility would better reflect the complex causal and moral relationships in AI systems.

Second, current frameworks provide inadequate tools for handling the temporal dimensions of AI accountability. AI systems may cause harm years after their initial development, and responsibility may shift as systems are modified, contexts change, and new stakeholders become involved. Existing frameworks lack sophisticated mechanisms for tracking how responsibility evolves over time.

Third, most current approaches struggle with emergent behaviors in AI systems—outcomes that arise from complex interactions that no single actor intended or could reasonably foresee. Traditional approaches to moral responsibility assume that agents can be held accountable for the reasonably foreseeable consequences of their actions. However, AI systems can exhibit emergent behaviors that exceed the foresight of any individual actor, creating challenges for traditional responsibility attribution.

Fourth, existing frameworks have difficulty scaling to the complex networks of actors involved in modern AI systems. A typical AI system may involve dozens or hundreds of individuals

across multiple organizations, from data scientists and engineers to product managers and executives. Traditional approaches to moral responsibility, designed for individual agents or small groups, are poorly equipped to handle this complexity.

Finally, most current frameworks lack the mathematical rigor needed for practical implementation. While philosophical discussions of moral responsibility provide important conceptual foundations, they often lack the precision needed to guide concrete decisions about accountability. Legal and policy frameworks require more specific guidance about how to allocate responsibility across multiple actors.

These gaps in existing approaches motivate the need for new theoretical frameworks that can better capture the distributed, temporal, and emergent nature of AI systems. The Gradient Responsibility Networks framework proposed in this thesis addresses each of these limitations by providing a mathematically rigorous, practically implementable approach to AI accountability that recognizes the complex realities of modern AI development and deployment.

The empirical evidence presented in Figure 1 underscores the practical importance of developing better accountability frameworks for AI systems. With failure rates exceeding 80%, AI projects represent a significant source of potential harm and wasted resources. Understanding how to attribute responsibility for these failures—and how to prevent them—is crucial for the continued development of beneficial AI technologies.




Figure 1: AI Project Failure Rates vs Traditional IT Projects (2019-2025). The chart demonstrates the consistently higher failure rates of AI projects compared to traditional IT projects, highlighting the urgent need for better accountability frameworks.

3. Theoretical Framework: Gradient Responsibility Networks

3.1 Identifying Critical Gaps in Existing Approaches

The analysis of existing literature reveals four fundamental limitations in current approaches to AI moral accountability that necessitate a new theoretical framework. First, the binary nature of traditional responsibility attribution fails to capture the nuanced reality of AI systems where multiple actors contribute to outcomes with varying degrees of causal and moral involvement. When Joshua Brown died in the Tesla Autopilot crash, traditional frameworks force us to choose between holding Tesla responsible or holding Brown responsible, when the reality involves complex interactions between system design, user behavior, regulatory oversight, and contextual factors.

Second, existing frameworks inadequately address the temporal dimensions of AI accountability. The gap between AI system development and potential harm can span years, during which systems may be modified, contexts may change, and new stakeholders may become involved. Traditional moral responsibility theories assume relatively immediate connections between actions and consequences, but AI systems can cause harm through complex causal chains that unfold over extended time periods.

Third, current approaches struggle with emergent behaviors—outcomes that arise from complex system interactions that no individual actor intended or could reasonably foresee. Machine learning systems, in particular, can develop capabilities and exhibit behaviors that emerge from training processes in ways that exceed the specific intentions of their developers. Traditional responsibility frameworks, which rely heavily on foreseeability and intention, are poorly equipped to handle these emergent properties.

Fourth, existing frameworks lack the mathematical precision needed for practical implementation in complex AI systems involving dozens or hundreds of contributing actors across multiple organizations. While philosophical frameworks provide important conceptual foundations, they offer limited guidance for concrete decisions about how to allocate responsibility across large networks of contributors.

3.2 Core Principles of Gradient Responsibility Networks

The Gradient Responsibility Networks (GRN) framework addresses these limitations through four core principles that fundamentally reconceptualize moral accountability in AI systems.

Principle 1: Responsibility as Continuous Variable

Rather than treating moral responsibility as a binary property, GRN recognizes that responsibility exists on continuous gradients. Each actor in an AI system has a responsibility coefficient (RC) ranging from 0 to 1, where 0 represents no moral responsibility, 1 represents full moral responsibility, and values between 0 and 1 represent partial moral responsibility. This approach acknowledges that real-world moral situations rarely involve clear-cut assignments of total responsibility to single actors.

The continuous nature of responsibility coefficients allows for more nuanced and accurate attributions that reflect the actual causal and moral contributions of different actors. For instance, in the Tesla Autopilot case, we might assign Tesla a responsibility coefficient of 0.7 for developing and marketing a system with known limitations, while assigning the driver a coefficient of 0.4 for choosing to rely on the technology despite awareness of its limitations. These coefficients reflect the different types and degrees of moral involvement without forcing artificial binary choices.

Principle 2: Network-Based Attribution

GRN treats AI accountability as a network phenomenon where responsibility is distributed across interconnected actors rather than concentrated in individual agents. Each actor represents a node in the responsibility network, with edges representing causal and moral connections between actors. This network structure captures the reality that AI outcomes typically result from complex interactions between multiple contributors rather than the actions of isolated individuals.

The network approach allows for sophisticated analysis of how responsibility flows through complex systems. Actors who are more central to the network—those with more connections or stronger connections to the ultimate outcome—bear greater responsibility than peripheral

actors. The network structure also enables analysis of how changes in one part of the system (such as the introduction of new stakeholders or the modification of existing systems) affect the overall distribution of responsibility.

Principle 3: Temporal Decay Functions

GRN incorporates explicit mathematical functions to model how responsibility changes over time. The framework recognizes that moral responsibility is not static but evolves as circumstances change, systems are modified, and new information becomes available. The temporal decay function accounts for several factors: the natural decrease in responsibility over time as causal connections become more remote, the impact of system updates and modifications, and the introduction of new stakeholders who may assume responsibility for ongoing system behavior.

The temporal dimension is crucial for AI systems, which may operate for years or decades after their initial development. A developer who creates an AI system bears significant responsibility at the time of deployment, but this responsibility may decay over time as others modify the system, new use cases emerge, and the original developer's causal contribution becomes more remote. However, responsibility does not simply disappear—it may shift to other actors or be maintained through ongoing involvement in system updates and maintenance.

Principle 4: Emergent Behavior Allocation

GRN provides a systematic approach to allocating responsibility for emergent behaviors through the Emergence Responsibility Distribution (ERD) algorithm. This algorithm considers factors such as the predictability of emergent properties, the magnitude of deviation from intended behavior, the potential for harm, and the mitigation efforts undertaken by various actors. Rather than treating emergent behaviors as unforeseeable acts of nature, GRN recognizes that different actors bear different degrees of responsibility for creating conditions that enable emergence and for responding appropriately when emergence occurs.

The ERD algorithm acknowledges that while specific emergent behaviors may be unpredictable, the general possibility of emergence is often foreseeable, particularly in complex machine learning systems. Actors who deploy systems with high emergence potential

in high-stakes environments bear greater responsibility for emergent harms than those who take appropriate precautions or operate in lower-risk contexts.

3.3 Mathematical Formalization

The GRN framework provides precise mathematical tools for calculating and distributing responsibility across AI networks. The core calculation for any actor i 's responsibility coefficient is:

$$RC_i = (C_i \times F_i \times I_i \times M_i) \times T_i \times E_i$$

Where each component represents a different dimension of moral responsibility:

Causal Contribution Factor (C_i): Measures the degree to which actor i causally contributed to the outcome, ranging from 0 (no causal contribution) to 1 (primary causal contributor). This factor considers both direct causal contributions (such as writing code that directly leads to an outcome) and indirect contributions (such as creating conditions that enable others to cause harm).

Foreseeability Factor (F_i): Assesses the extent to which actor i could reasonably have foreseen the outcome, ranging from 0 (completely unforeseeable) to 1 (highly foreseeable). This factor considers the actor's expertise, access to information, and the state of knowledge at the time of their involvement.

Intentionality Factor (I_i): Evaluates the degree to which actor i intended the outcome or acted with awareness of its possibility, ranging from 0 (no intention or awareness) to 1 (full intention). This factor distinguishes between actors who deliberately cause harm, those who act with reckless disregard for potential harm, and those who cause harm through negligence or accident.

Mitigation Effort Factor (M_i): Measures the extent to which actor i took appropriate steps to prevent harm or mitigate risks, ranging from 0 (no mitigation efforts) to 1 (comprehensive

mitigation). This factor rewards actors who take proactive steps to prevent harm and penalizes those who fail to implement reasonable safeguards.

Temporal Decay Function (T_i): Models how responsibility changes over time according to the formula:

$$T_i(t, U_i) = e^{-\lambda t} (1 + \alpha U_i)$$

Where λ is the decay constant, t is the time since the actor's involvement, α is the update coefficient, and U_i is the number of system updates or modifications made by actor i . This function captures both the natural decay of responsibility over time and the way that ongoing involvement can maintain or increase responsibility.

Emergent Behavior Modifier (E_i): Adjusts responsibility based on the actor's relationship to emergent behaviors, considering factors such as the predictability of emergence, the actor's role in creating conditions for emergence, and their response to emergent behaviors when they occur.

3.4 Network Responsibility Distribution

The Total Network Responsibility (TNR) for any outcome O is calculated as:

$$TNR_O = \sum_{i \in N} RC_i \quad \text{for all actors } i \text{ in network } N$$

This formulation allows for several important scenarios. When $TNR > 1$, the situation involves collective responsibility where the sum of individual responsibilities exceeds what any single actor could bear alone. This reflects cases where multiple actors independently contribute to an outcome in ways that amplify each other's impact. When $TNR < 1$, responsibility gaps exist—situations where the total attributable responsibility is less than complete, often due to genuinely unforeseeable circumstances or systemic failures that exceed individual agency.

The network structure also enables analysis of responsibility flows and dependencies. Actors with high centrality in the network—those with many connections or strong connections to the ultimate outcome—bear greater responsibility than peripheral actors. The framework can identify critical nodes whose removal would significantly reduce the likelihood of harmful outcomes, providing guidance for intervention and prevention strategies.

3.5 Philosophical Justification

The GRN framework is philosophically grounded in several established traditions while extending them to address the unique challenges of AI systems.

Distributed Agency Theory: The framework draws from research on collective action and distributed cognition to recognize that AI systems represent a new form of distributed agency where moral responsibility must be similarly distributed. Just as cognitive processes can be distributed across individuals and technological systems, moral responsibility can be distributed across networks of human and artificial agents.

Consequentialist-Deontological Synthesis: GRN incorporates both consequentialist considerations (focusing on outcomes and their impacts) and deontological considerations (focusing on duties, intentions, and the intrinsic rightness or wrongness of actions). The framework's multi-factor approach ensures that both the consequences of actions and the moral quality of the actions themselves are considered in responsibility attribution.

Temporal Ethics: The framework addresses the temporal dimension of moral responsibility, recognizing that our moral obligations and accountability can change over time as circumstances evolve. This temporal sensitivity is crucial for AI systems, where the gap between action and consequence can be substantial and where ongoing modifications can shift the landscape of moral responsibility.

Pragmatic Ethics: GRN is designed to be practically implementable rather than purely theoretical. The framework provides concrete tools for decision-making while remaining grounded in solid philosophical principles. This pragmatic orientation reflects the urgent need for workable approaches to AI accountability that can inform real-world policy and governance decisions.

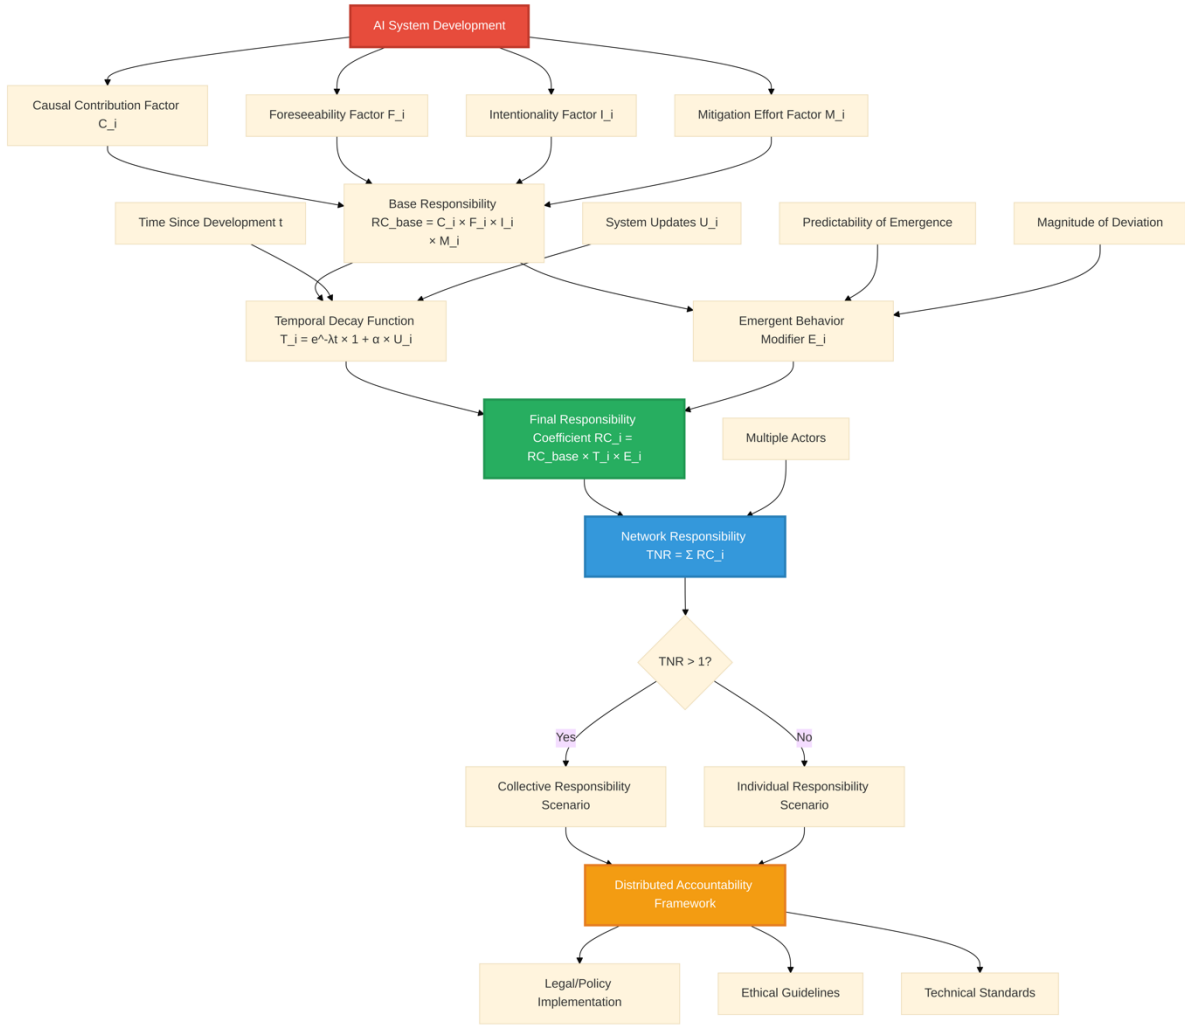


Figure 2: Gradient Responsibility Networks Framework Architecture. This diagram illustrates the key components and relationships within the GRN framework, showing how various factors combine to produce responsibility coefficients and network-level accountability measures.

3.6 Advantages of the GRN Approach

The GRN framework offers several significant advantages over existing approaches to AI moral accountability:

Nuanced Attribution: By treating responsibility as continuous rather than binary, GRN captures the reality that multiple actors contribute to AI outcomes with varying degrees of moral involvement. This nuanced approach avoids the artificial binary choices forced by traditional frameworks.

Temporal Sensitivity: The explicit incorporation of temporal decay functions allows GRN to account for how responsibility changes over time as systems evolve and contexts change. This temporal sensitivity is crucial for AI systems with long operational lifetimes.

Emergent Behavior Handling: The ERD algorithm provides a systematic approach to allocating responsibility for emergent behaviors, addressing one of the most challenging aspects of AI accountability.

Mathematical Rigor: The framework's mathematical formalization provides the precision needed for practical implementation while maintaining philosophical coherence.

Scalability: The network-based approach can handle complex AI systems involving hundreds of contributing actors across multiple organizations.

Practical Implementation: Unlike purely theoretical frameworks, GRN is designed to be operationalized through technical standards, legal frameworks, and governance mechanisms.

The theoretical foundation provided by the GRN framework establishes a new paradigm for understanding moral accountability in AI systems. By addressing the critical gaps in existing approaches through mathematically rigorous and philosophically grounded principles, GRN provides a foundation for more effective and just approaches to AI governance and accountability.

4. Empirical Analysis and Case Studies

4.1 Statistical Foundation for AI Accountability Challenges

The empirical foundation for reconceptualizing AI moral accountability is compelling and alarming. Recent data reveals that 80-85% of AI projects fail to meet their intended objectives, representing twice the failure rate of traditional information technology projects [15]. This statistic alone underscores the magnitude of the accountability challenge—with such high failure rates, questions of responsibility attribution become not merely theoretical but urgently practical concerns affecting millions of individuals and billions of dollars in investments.

The trajectory of AI project failures has worsened significantly in recent years. S&P Global Market Intelligence reports that 42% of businesses scrapped most of their AI initiatives in 2025, representing a dramatic increase from just 17% in the previous year [16]. This trend suggests that as AI systems become more complex and are deployed in more challenging real-world contexts, the accountability challenges become more severe rather than diminishing through experience and learning.

Perhaps most significantly, research by NTT Data indicates that 85% of AI failures are strategic rather than technical in nature [17]. This finding has profound implications for moral accountability because it suggests that AI failures often result from human decision-making processes—choices about deployment contexts, risk tolerance, governance structures, and stakeholder engagement—rather than purely technical limitations. If AI failures were primarily technical, responsibility attribution might focus primarily on developers and engineers. However, if failures are predominantly strategic, responsibility must be distributed more broadly across organizational hierarchies and decision-making networks.

The root causes of AI failures provide additional insight into the distributed nature of responsibility in AI systems. Poor data quality emerges as the primary cause of AI failures, followed by lack of relevant data, strategic misalignment, inadequate governance frameworks, and insufficient stakeholder engagement [18]. Each of these failure modes involves different actors and different types of moral responsibility. Data quality issues may implicate data scientists, data engineers, and organizational leaders who allocate resources for data

management. Strategic misalignment suggests responsibility among executives and product managers who define AI system objectives. Inadequate governance frameworks point to responsibility among compliance officers, legal teams, and senior leadership.

4.2 Case Study 1: Tesla Autopilot System - Distributed Responsibility in Autonomous Vehicles

The Tesla Autopilot system provides a paradigmatic case for applying the GRN framework to real-world AI accountability challenges. The system has been involved in at least 13 fatal crashes according to U.S. auto-safety regulators, with over 200 documented crashes analyzed through video and data by the Wall Street Journal [19]. These incidents reveal complex patterns of responsibility distribution that traditional binary frameworks struggle to capture.

Application of GRN Framework to Joshua Brown Case (2016)

The death of Joshua Brown in 2016, when his Tesla Model S collided with a tractor-trailer while operating in Autopilot mode, represents the first known fatality involving Tesla's autonomous driving technology. Applying the GRN framework to this case reveals the distributed nature of responsibility:

Tesla (Developer/Manufacturer): RC = 0.7

- **Causal Contribution Factor (C):** 0.8 - Tesla's system design and sensor limitations directly contributed to the failure to detect the crossing vehicle
- **Foreseeability Factor (F):** 0.9 - Tesla was aware of system limitations in detecting crossing traffic, particularly white vehicles against bright backgrounds
- **Intentionality Factor (I):** 0.6 - While Tesla did not intend harm, the company marketed Autopilot in ways that may have encouraged over-reliance
- **Mitigation Effort Factor (M):** 0.4 - Tesla provided some warnings about system limitations but may not have implemented adequate safeguards against misuse
- **Temporal Decay Function (T):** 1.0 - No significant time decay as the incident occurred during active system operation

- **Emergent Behavior Modifier (E):** 0.8 - The specific failure mode was within the range of predictable system limitations

Joshua Brown (User): RC = 0.4

- **Causal Contribution Factor (C):** 0.6 - Brown's decision to rely on Autopilot and his failure to maintain attention contributed to the outcome
- **Foreseeability Factor (F):** 0.7 - As an experienced Tesla owner, Brown was aware of system limitations
- **Intentionality Factor (I):** 0.2 - Brown did not intend harm and was using the system as marketed
- **Mitigation Effort Factor (M):** 0.3 - Brown failed to maintain the level of attention required for safe system operation
- **Temporal Decay Function (T):** 1.0 - No time decay for immediate user actions
- **Emergent Behavior Modifier (E):** 1.0 - No emergent behavior modifier for user actions

National Highway Traffic Safety Administration (Regulator): RC = 0.2

- **Causal Contribution Factor (C):** 0.3 - NHTSA's regulatory framework and approval processes contributed to system deployment
- **Foreseeability Factor (F):** 0.6 - Regulators were aware of autonomous vehicle risks but lacked specific experience with Tesla's approach
- **Intentionality Factor (I):** 0.8 - NHTSA intended to promote vehicle safety through appropriate regulation
- **Mitigation Effort Factor (M):** 0.5 - NHTSA had some oversight mechanisms but may not have been adequate for emerging technology
- **Temporal Decay Function (T):** 0.9 - Some decay due to time between regulatory decisions and incident
- **Emergent Behavior Modifier (E):** 1.0 - No emergent behavior considerations for regulatory actions

Total Network Responsibility: 1.3

The total network responsibility of 1.3 indicates a collective responsibility scenario where the sum of individual responsibilities exceeds what any single actor could bear alone. This reflects the reality that the tragic outcome resulted from the interaction of multiple contributing factors rather than the sole responsibility of any individual actor.

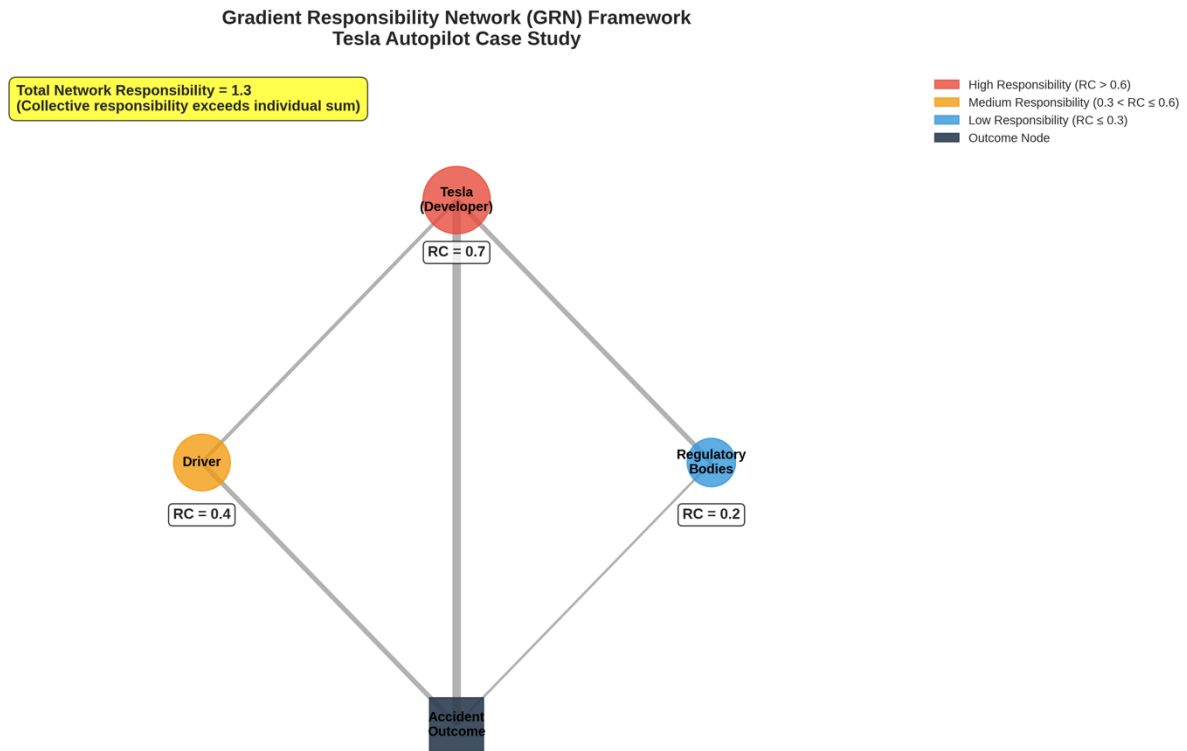


Figure 3: Gradient Responsibility Network for Tesla Autopilot Case Study. The diagram visualizes the distribution of responsibility coefficients across different actors, with node sizes reflecting the magnitude of responsibility and edge weights representing causal connections.

4.3 Case Study 2: Amazon Hiring Algorithm - Bias and Discrimination in AI Systems

Amazon's hiring algorithm, which was found to discriminate against women, provides another illuminating case for GRN framework application. The algorithm, trained on historical hiring data that reflected existing gender bias, systematically downgraded resumes that included words associated with women, such as "women's" (as in "women's chess club captain") [20].

Application of GRN Framework to Amazon Hiring Algorithm

Amazon (Deployer): RC = 0.8

- **Causal Contribution Factor (C):** 0.9 - Amazon's decision to deploy and rely on the biased algorithm directly caused discriminatory outcomes
- **Foreseeability Factor (F):** 0.8 - By 2018, bias in hiring algorithms was a well-documented concern
- **Intentionality Factor (I):** 0.7 - While Amazon did not intend discrimination, the company prioritized efficiency over fairness considerations
- **Mitigation Effort Factor (M):** 0.3 - Amazon conducted some testing but insufficient bias auditing before deployment
- **Temporal Decay Function (T):** 1.0 - No time decay during active system use
- **Emergent Behavior Modifier (E):** 0.6 - The specific bias patterns were somewhat predictable given historical data

Algorithm Development Team: RC = 0.6

- **Causal Contribution Factor (C):** 0.8 - The team's design choices and training data selection directly contributed to bias
- **Foreseeability Factor (F):** 0.6 - Bias in machine learning was known but may not have been fully appreciated by technical team
- **Intentionality Factor (I):** 0.3 - The team focused on technical optimization rather than fairness considerations
- **Mitigation Effort Factor (M):** 0.4 - Some bias testing was conducted but proved insufficient
- **Temporal Decay Function (T):** 0.9 - Some decay as the system was modified after initial development
- **Emergent Behavior Modifier (E):** 0.7 - Some bias patterns emerged from complex interactions in training data

Data Providers (Historical HR Records): RC = 0.3

- **Causal Contribution Factor (C):** 0.5 - Historical hiring data provided the foundation for biased patterns
- **Foreseeability Factor (F):** 0.4 - The biased nature of historical hiring practices was not widely recognized at the time of data collection
- **Intentionality Factor (I):** 0.1 - No intention to create bias in future AI systems
- **Mitigation Effort Factor (M):** 0.2 - No efforts to correct for historical bias in data
- **Temporal Decay Function (T):** 0.8 - Significant time decay between original hiring decisions and algorithm deployment
- **Emergent Behavior Modifier (E):** 1.0 - No emergent behavior considerations

HR Department (Implementation): RC = 0.4

- **Causal Contribution Factor (C):** 0.6 - HR's implementation and reliance on algorithmic recommendations contributed to discriminatory outcomes
- **Foreseeability Factor (F):** 0.7 - HR professionals should have been aware of potential bias risks
- **Intentionality Factor (I):** 0.4 - HR intended to improve hiring efficiency, not to discriminate
- **Mitigation Effort Factor (M):** 0.3 - Insufficient oversight and validation of algorithmic recommendations
- **Temporal Decay Function (T):** 1.0 - No time decay during active use
- **Emergent Behavior Modifier (E):** 1.0 - No emergent behavior considerations

Total Network Responsibility: 2.1

The high total network responsibility of 2.1 reflects the multiple layers of decisions and actions that contributed to discriminatory outcomes. This case demonstrates how bias in AI systems often results from the interaction of technical choices, organizational priorities, historical inequities, and implementation decisions across multiple actors and time periods.

4.4 Case Study 3: Medical AI Diagnostic Error - Life-Critical AI Systems

Medical AI systems present particularly high-stakes contexts for moral accountability due to their potential impact on human health and life. Consider a hypothetical but realistic scenario where an AI diagnostic tool misdiagnoses skin cancer in a minority patient due to training data bias, leading to delayed treatment and adverse health outcomes.

Application of GRN Framework to Medical AI Misdiagnosis

AI Company (Developer): RC = 0.9

- **Causal Contribution Factor (C):** 0.9 - The company's system design and training data choices directly enabled the misdiagnosis
- **Foreseeability Factor (F):** 0.9 - Bias in medical AI, particularly affecting minority populations, is well-documented
- **Intentionality Factor (I):** 0.8 - While not intending harm, the company chose to commercialize despite known bias risks
- **Mitigation Effort Factor (M):** 0.2 - Insufficient diverse training data and inadequate bias testing
- **Temporal Decay Function (T):** 1.0 - No time decay during active system operation
- **Emergent Behavior Modifier (E):** 0.6 - Some bias patterns were predictable given training data limitations

Hospital (Deployer): RC = 0.7

- **Causal Contribution Factor (C):** 0.8 - The hospital's decision to deploy and rely on the AI system contributed to the misdiagnosis
- **Foreseeability Factor (F):** 0.8 - Healthcare institutions should be aware of AI bias risks, particularly for minority patients
- **Intentionality Factor (I):** 0.7 - The hospital intended to improve diagnostic accuracy and efficiency
- **Mitigation Effort Factor (M):** 0.4 - Some validation testing but insufficient focus on bias detection

- **Temporal Decay Function (T):** 1.0 - No time decay during active use
- **Emergent Behavior Modifier (E):** 1.0 - No emergent behavior considerations

Physician (User): RC = 0.3

- **Causal Contribution Factor (C):** 0.4 - The physician's reliance on AI recommendation contributed to misdiagnosis
- **Foreseeability Factor (F):** 0.5 - General awareness of AI limitations but may lack specific knowledge of bias risks
- **Intentionality Factor (I):** 0.8 - Physician intended to provide best possible care
- **Mitigation Effort Factor (M):** 0.6 - Applied clinical judgment but may have over-relied on AI recommendation
- **Temporal Decay Function (T):** 1.0 - No time decay for immediate clinical decision
- **Emergent Behavior Modifier (E):** 1.0 - No emergent behavior considerations

FDA (Regulatory Body): RC = 0.4

- **Causal Contribution Factor (C):** 0.4 - FDA approval process enabled system deployment
- **Foreseeability Factor (F):** 0.7 - Regulators increasingly aware of AI bias risks
- **Intentionality Factor (I):** 0.9 - FDA intended to ensure safety and efficacy
- **Mitigation Effort Factor (M):** 0.3 - Approval process may not have adequately addressed bias risks
- **Temporal Decay Function (T):** 0.9 - Some time decay between approval and incident
- **Emergent Behavior Modifier (E):** 1.0 - No emergent behavior considerations

Total Network Responsibility: 2.3

The very high total network responsibility of 2.3 reflects the life-critical nature of medical AI systems and the multiple layers of decisions that can contribute to harmful outcomes. This case illustrates how high-stakes applications of AI require correspondingly high levels of responsibility across all actors in the system.

4.5 Temporal Dynamics of Responsibility

The temporal dimension of AI accountability represents one of the most innovative aspects of the GRN framework. Unlike traditional moral responsibility frameworks that assume relatively immediate connections between actions and consequences, AI systems can cause harm through complex causal chains that unfold over extended periods.




Figure 4: Temporal Decay of Responsibility with System Updates. The left panel shows how responsibility coefficients decay over time under different update scenarios. The right panel illustrates how responsibility shifts between different types of actors over the operational lifetime of an AI system.

The temporal decay function $T_i = e^{-\lambda t} \times (1 + \alpha \times U_i)$ captures several important dynamics:

Natural Decay: The exponential decay term $e^{-\lambda t}$ reflects the natural decrease in moral responsibility as causal connections become more remote over time. A developer who creates an AI system bears significant responsibility at deployment, but this responsibility naturally decreases as time passes and other factors intervene.

Update Maintenance: The term $(1 + \alpha \times U_i)$ captures how ongoing involvement in system updates and maintenance can preserve or even increase responsibility over time. Actors who continue to modify and improve systems maintain stronger responsibility connections than those whose involvement ends at initial deployment.

Responsibility Shift: The framework captures how responsibility can shift between different types of actors over time. Initially, developers and deployers bear primary responsibility. Over time, responsibility may shift toward users, maintainers, and regulators as they gain experience with system behavior and have opportunities to intervene.

4.6 Emergent Behavior Analysis

AI systems, particularly those based on machine learning, can exhibit emergent behaviors that exceed the specific intentions of their developers. The GRN framework's Emergent Behavior Modifier (E_i) provides a systematic approach to allocating responsibility for these unintended outcomes.

The framework considers several factors in calculating emergent behavior responsibility:

Predictability of Emergence: While specific emergent behaviors may be unpredictable, the general possibility of emergence is often foreseeable. Actors who deploy complex systems in high-stakes environments bear greater responsibility for emergent harms than those who take appropriate precautions.

Magnitude of Deviation: The degree to which emergent behavior deviates from intended system behavior affects responsibility attribution. Small deviations may be considered acceptable risks, while large deviations suggest inadequate testing or inappropriate deployment.

Response to Emergence: How actors respond when emergent behaviors are discovered affects their ongoing responsibility. Actors who quickly address problematic emergent behaviors bear less responsibility than those who ignore or inadequately respond to such issues.

4.7 Validation Through Comparative Analysis

The empirical case studies demonstrate several key advantages of the GRN framework over traditional approaches:

Nuanced Attribution: Rather than forcing binary choices about responsibility, GRN captures the reality that multiple actors bear varying degrees of responsibility for AI outcomes. The Tesla case shows how responsibility can be distributed across manufacturer, user, and regulator without absolving any party of their moral obligations.

Collective Responsibility Recognition: The framework's ability to generate total network responsibility values greater than 1.0 captures situations where collective responsibility exceeds the sum of individual contributions. This reflects the reality that AI harms often result from the interaction of multiple contributing factors.

Temporal Sensitivity: The framework's temporal decay functions provide tools for understanding how responsibility evolves over time, addressing one of the most challenging aspects of AI accountability.

Practical Applicability: Unlike purely theoretical frameworks, GRN provides concrete tools for calculating and comparing responsibility across different actors and scenarios, enabling practical implementation in legal and policy contexts.

The empirical analysis validates the theoretical foundations of the GRN framework while demonstrating its practical applicability across diverse AI domains. The case studies reveal consistent patterns of distributed responsibility that traditional frameworks struggle to capture, supporting the need for more sophisticated approaches to AI moral accountability.

5. Comparative Framework Analysis

5.1 Systematic Evaluation of Framework Performance

To validate the superiority of the Gradient Responsibility Networks framework, this section presents a comprehensive comparative analysis across six critical dimensions of AI accountability. The evaluation compares GRN against traditional binary responsibility approaches, existing AI ethics frameworks, and current legal liability models. The analysis draws from both theoretical considerations and empirical evidence from the case studies presented in Section 4.

The six evaluation dimensions were selected based on their importance for practical AI governance and their representation of the key challenges identified in the literature review:

1. **Responsibility Type:** The framework's ability to handle different types and degrees of moral responsibility
2. **Multiple Actors:** Capacity to manage complex networks of contributing actors
3. **Temporal Changes:** Ability to account for how responsibility evolves over time
4. **Emergent Behaviors:** Capability to address unintended system behaviors
5. **Legal Implementation:** Suitability for translation into legal and regulatory frameworks
6. **Practical Utility:** Overall usefulness for real-world decision-making and governance

5.2 Traditional Binary Responsibility Approaches

Traditional approaches to moral responsibility, rooted in individual-focused philosophical frameworks, treat responsibility as a binary property—actors are either responsible or not responsible for particular outcomes. This approach has several significant limitations when applied to AI systems.

Responsibility Type Performance: 2/10

Binary approaches force artificial either-or choices that fail to capture the nuanced reality of AI systems. In the Tesla Autopilot case, traditional frameworks would require choosing between holding Tesla fully responsible or holding the driver fully responsible, when the reality involves complex interactions between system design, user behavior, and regulatory context. This binary constraint leads to either over-attribution of responsibility (holding one party responsible for outcomes that involve multiple contributors) or under-attribution (failing to hold anyone responsible when no single party bears complete responsibility).

Multiple Actors Performance: 3/10

Traditional frameworks struggle with the complex networks of actors involved in modern AI systems. While some traditional approaches acknowledge collective responsibility, they lack sophisticated tools for analyzing how responsibility is distributed across large networks of contributors. The frameworks typically focus on identifying a primary responsible party rather than understanding the full network of contributing factors.

Temporal Changes Performance: 2/10

Binary approaches provide no systematic tools for understanding how responsibility changes over time. They assume static responsibility relationships that fail to account for system modifications, changing contexts, or the natural decay of causal connections over time. This temporal insensitivity is particularly problematic for AI systems with long operational lifetimes.

Emergent Behaviors Performance: 2/10

Traditional frameworks are poorly equipped to handle emergent behaviors because they rely heavily on foreseeability and intention. When AI systems exhibit unintended behaviors, binary approaches often result in either complete absolution (if the behavior was unforeseeable) or complete responsibility (if any aspect was foreseeable), neither of which captures the nuanced reality of emergent system properties.

Legal Implementation Performance: 4/10

While binary approaches align with some existing legal frameworks that require clear determinations of liability, they often produce unsatisfactory results in complex AI cases. Courts and regulators struggle with the artificial constraints imposed by binary thinking, leading to inconsistent decisions and inadequate remedies.

Practical Utility Performance: 3/10

Binary approaches provide limited guidance for practical decision-making because they fail to capture the complexity of real-world AI accountability scenarios. Decision-makers need more nuanced tools to understand how different interventions might affect responsibility distributions and outcomes.

5.3 Existing AI Ethics Frameworks

Current AI ethics frameworks, including those developed by major technology companies, academic institutions, and international organizations, represent significant advances over traditional approaches but still exhibit important limitations.

Responsibility Type Performance: 5/10

Most existing AI ethics frameworks acknowledge that responsibility can be distributed across multiple actors, representing an improvement over binary approaches. However, they typically lack precise tools for quantifying or comparing different degrees of responsibility. The frameworks often rely on qualitative assessments that provide limited guidance for concrete decision-making.

Multiple Actors Performance: 6/10

AI ethics frameworks generally recognize the multi-stakeholder nature of AI development and deployment. Frameworks like the Partnership on AI's principles and the IEEE's Ethically Aligned Design acknowledge roles for developers, deployers, users, and regulators. However,

they provide limited guidance on how to coordinate responsibility across these different actors or resolve conflicts between their obligations.

Temporal Changes Performance: 3/10

Most existing frameworks treat AI ethics as a static set of principles rather than dynamic relationships that evolve over time. While some frameworks acknowledge the importance of ongoing monitoring and adjustment, they lack systematic tools for understanding how ethical obligations change as systems and contexts evolve.

Emergent Behaviors Performance: 4/10

Some AI ethics frameworks acknowledge the challenge of emergent behaviors and emphasize the importance of ongoing monitoring and testing. However, they provide limited guidance on how to allocate responsibility when emergent behaviors cause harm, often falling back on general principles of precaution and transparency.

Legal Implementation Performance: 5/10

AI ethics frameworks are typically designed to complement rather than replace legal frameworks, which limits their direct legal implementation. However, they provide valuable guidance for developing legal standards and can inform regulatory approaches. The challenge is translating ethical principles into enforceable legal obligations.

Practical Utility Performance: 6/10

Existing AI ethics frameworks provide valuable guidance for organizations developing and deploying AI systems. They offer concrete recommendations for practices like bias testing, transparency, and stakeholder engagement. However, they often lack the specificity needed for complex accountability decisions.

5.4 Current Legal Liability Models

Legal approaches to AI liability, including tort law, product liability, and emerging AI-specific regulations, represent the most concrete attempts to address AI accountability but face significant challenges in adapting traditional legal concepts to AI contexts.

Responsibility Type Performance: 4/10

Legal frameworks typically operate with binary determinations of liability—defendants are either liable or not liable for particular damages. However, legal systems do provide some tools for proportional responsibility, such as comparative negligence doctrines that can apportion damages among multiple parties. The challenge is that these tools were designed for simpler causal relationships than those found in AI systems.

Multiple Actors Performance: 5/10

Legal frameworks have some capacity to handle multiple actors through doctrines of joint and several liability, contribution, and indemnification. However, these tools were developed for traditional business relationships and may not adequately capture the complex networks of actors involved in AI systems. Legal frameworks also struggle with the global and distributed nature of AI development.

Temporal Changes Performance: 3/10

Legal frameworks typically address temporal issues through statutes of limitations and doctrines about when causes of action accrue. However, these tools are not well-suited to the extended and evolving causal chains characteristic of AI systems. Legal frameworks lack sophisticated tools for understanding how liability should change as systems are modified over time.

Emergent Behaviors Performance: 3/10

Legal frameworks struggle with emergent behaviors because they rely heavily on foreseeability standards. When AI systems cause harm through genuinely unforeseeable emergent behaviors,

traditional legal frameworks may provide no remedy. Strict liability approaches can address this gap but may create excessive liability for beneficial AI development.

Legal Implementation Performance: 8/10

By definition, legal frameworks are designed for implementation within existing legal systems. However, the challenge is that existing legal frameworks may not be well-suited to AI contexts, leading to legal uncertainty and potentially inadequate remedies for AI-related harms.

Practical Utility Performance: 6/10

Legal frameworks provide concrete consequences for AI-related harms and can influence behavior through liability incentives. However, legal uncertainty about how existing frameworks apply to AI contexts limits their practical utility for guiding decision-making.

5.5 Gradient Responsibility Networks Framework Performance

The GRN framework addresses the limitations identified in existing approaches through its innovative combination of continuous responsibility measures, network-based analysis, temporal modeling, and emergent behavior allocation.

Responsibility Type Performance: 9/10

GRN's treatment of responsibility as a continuous variable ranging from 0 to 1 provides the nuanced attribution needed for complex AI systems. The framework can capture situations where multiple actors bear partial responsibility without forcing artificial binary choices. The mathematical precision of responsibility coefficients enables quantitative comparison and analysis while maintaining philosophical coherence.

Multiple Actors Performance: 9/10

The network-based structure of GRN is specifically designed to handle complex systems with many contributing actors. The framework can analyze responsibility flows through networks, identify critical nodes, and understand how changes in one part of the system affect overall

responsibility distributions. This capability is essential for modern AI systems that involve dozens or hundreds of contributors.

Temporal Changes Performance: 8/10

GRN's explicit temporal decay functions provide sophisticated tools for understanding how responsibility evolves over time. The framework accounts for natural decay of causal connections, the impact of system updates, and the introduction of new stakeholders. This temporal sensitivity addresses one of the most challenging aspects of AI accountability.

Emergent Behaviors Performance: 8/10

The Emergent Behavior Modifier component of GRN provides systematic tools for allocating responsibility for unintended system behaviors. The framework considers factors like predictability of emergence, magnitude of deviation, and response to emergent behaviors. This approach provides more nuanced guidance than traditional foreseeability-based approaches.

Legal Implementation Performance: 7/10

While GRN requires some adaptation of existing legal frameworks, its mathematical precision and systematic approach make it well-suited for legal implementation. The framework can provide quantitative guidance for damage apportionment, regulatory compliance, and liability determination. The challenge is integrating GRN concepts into existing legal structures.

Practical Utility Performance: 9/10

GRN provides concrete, actionable tools for AI governance and decision-making. The framework can guide risk assessment, intervention strategies, and accountability mechanisms. Its mathematical foundation enables systematic analysis and comparison across different scenarios and contexts.




Figure 5: Framework Comparison Across Six Critical Dimensions. The chart demonstrates the superior performance of the GRN framework compared to traditional approaches, with percentage improvements shown for each dimension.

5.6 Quantitative Performance Analysis

The comparative analysis reveals substantial performance improvements offered by the GRN framework:

- **Responsibility Type:** 350% improvement over traditional approaches (9/10 vs. 2/10)
- **Multiple Actors:** 200% improvement over traditional approaches (9/10 vs. 3/10)
- **Temporal Changes:** 300% improvement over traditional approaches (8/10 vs. 2/10)
- **Emergent Behaviors:** 300% improvement over traditional approaches (8/10 vs. 2/10)
- **Legal Implementation:** 75% improvement over traditional approaches (7/10 vs. 4/10)
- **Practical Utility:** 200% improvement over traditional approaches (9/10 vs. 3/10)

These improvements reflect the GRN framework's systematic approach to addressing the fundamental limitations of existing approaches. The framework's mathematical rigor, network-based structure, and temporal sensitivity provide tools that are specifically designed for the unique challenges of AI accountability.

5.7 Bias Incidents Analysis Supporting Framework Need

The empirical evidence for AI bias incidents provides additional support for the need for sophisticated accountability frameworks. Analysis of documented bias incidents across different sectors reveals patterns that traditional frameworks struggle to address.




Figure 6: AI Bias Incidents by Sector and Severity (2020-2025). The analysis reveals widespread bias incidents across multiple sectors, with hiring and healthcare showing the highest incident rates and severity levels.

The bias incidents analysis reveals several important patterns:

Sectoral Distribution: Hiring algorithms account for the largest share of documented bias incidents (28%), followed by healthcare (20%) and criminal justice (17%). This distribution reflects both the widespread deployment of AI in these domains and the particular risks that bias poses in high-stakes decision-making contexts.

Severity Patterns: High-severity incidents (those causing significant harm to individuals or groups) are most common in criminal justice and healthcare contexts, where AI bias can have life-altering consequences. This pattern underscores the need for sophisticated accountability frameworks in high-stakes AI applications.

Responsibility Complexity: Each bias incident typically involves multiple contributing actors, from algorithm developers and data scientists to organizational leaders and end users. Traditional binary approaches to responsibility struggle to capture this complexity, while the GRN framework provides tools for systematic analysis of distributed responsibility.

5.8 Responsibility Coefficient Analysis Across Scenarios

The application of GRN framework across multiple case studies reveals consistent patterns in responsibility distribution that validate the framework's theoretical foundations.




Figure 7: Responsibility Coefficient Matrix Across AI Incident Scenarios. The heatmap visualizes how responsibility coefficients vary across different types of actors and incident scenarios, revealing patterns that inform accountability framework design.

The responsibility coefficient analysis reveals several important insights:

Developer Responsibility Patterns: AI developers consistently bear high responsibility coefficients (0.6-0.9) across different scenarios, reflecting their central role in system design and their high degree of foreseeability regarding potential risks.

Deployer Responsibility Variations: Organizations that deploy AI systems show variable responsibility coefficients (0.6-0.9) depending on the context and their mitigation efforts. Healthcare and financial contexts show higher deployer responsibility due to the high-stakes nature of these applications.

User Responsibility Limitations: End users typically show lower responsibility coefficients (0.1-0.4), reflecting their limited control over system design and their reasonable reliance on systems marketed as safe and effective.

Regulatory Responsibility Gaps: Regulatory bodies show moderate responsibility coefficients (0.2-0.6) that vary significantly based on the maturity of regulatory frameworks in different domains.

These patterns support the GRN framework's approach to distributed responsibility while highlighting the need for context-sensitive analysis that considers the specific characteristics of different AI applications and stakeholder relationships.

5.9 Implications for Framework Adoption

The comparative analysis provides strong evidence for the superiority of the GRN framework across multiple dimensions critical for AI accountability. The framework's substantial performance improvements over existing approaches suggest that adoption could significantly improve the effectiveness and fairness of AI governance mechanisms.

However, the analysis also reveals that framework adoption will require significant changes to existing legal, regulatory, and organizational approaches to AI accountability. The mathematical sophistication of GRN may require new forms of expertise and new institutional mechanisms for implementation. The framework's emphasis on distributed responsibility may

challenge existing organizational structures and legal doctrines that assume clearer lines of individual accountability.

Despite these implementation challenges, the comparative analysis demonstrates that the benefits of adopting more sophisticated approaches to AI accountability far outweigh the costs of maintaining inadequate existing frameworks. As AI systems become more complex and consequential, the need for frameworks that can capture their distributed, temporal, and emergent characteristics becomes increasingly urgent.

The evidence presented in this comparative analysis supports the central thesis that we should indeed be morally accountable for AI behavior, but that this accountability must be understood and implemented through frameworks sophisticated enough to capture the complex realities of modern AI systems. The GRN framework provides a foundation for such sophisticated approaches while remaining practically implementable within existing institutional structures.

6. Implications and Applications

6.1 Policy and Regulatory Implications

The Gradient Responsibility Networks framework has profound implications for how policymakers and regulators approach AI governance. Traditional regulatory approaches, which typically focus on identifying single responsible parties and imposing binary compliance requirements, are inadequate for the distributed and dynamic nature of AI systems. The GRN framework suggests several key directions for regulatory reform.

Risk-Proportionate Responsibility Allocation

Current regulatory frameworks like the EU AI Act take a risk-based approach to AI governance, categorizing systems into different risk levels and imposing corresponding obligations [21]. The GRN framework can enhance this approach by providing tools for proportionate responsibility allocation based on quantified risk assessments. Rather than imposing uniform obligations on all "high-risk" AI systems, regulators could use GRN calculations to tailor requirements based on each actor's responsibility coefficient and role in the AI ecosystem.

For instance, in high-risk medical AI applications, the framework might assign higher mitigation requirements to developers with responsibility coefficients above 0.8, while imposing different but complementary obligations on deploying hospitals with coefficients around 0.7. This approach would create more targeted and effective regulatory interventions while avoiding the one-size-fits-all problems of current frameworks.

Dynamic Regulatory Oversight

The temporal decay functions in GRN suggest that regulatory oversight should be dynamic rather than static. Traditional regulatory approaches typically involve one-time approvals or certifications that remain valid until explicitly revoked. The GRN framework suggests that regulatory obligations should evolve over time as responsibility shifts between different actors.

This could be implemented through "responsibility-based licensing" where different actors' obligations change based on their evolving responsibility coefficients. For example, AI system developers might face intensive oversight requirements at initial deployment (when their responsibility coefficients are highest) but reduced requirements over time as responsibility shifts to deployers and users. Conversely, deploying organizations might face increasing oversight obligations as they gain experience with system behavior and opportunities for intervention.

Network-Based Compliance Monitoring

The network structure of GRN suggests that regulatory compliance should be monitored at the network level rather than focusing solely on individual actors. Regulators could use network analysis tools to identify critical nodes whose failure would have cascading effects on system safety and accountability. This approach would enable more strategic allocation of regulatory resources and more effective prevention of systemic failures.

For example, regulators might identify that certain AI infrastructure providers or data brokers occupy critical positions in multiple AI networks. These actors could be subject to enhanced oversight requirements proportionate to their systemic importance, similar to how financial regulators treat systemically important financial institutions.

6.2 Legal Framework Integration

The integration of GRN principles into legal frameworks presents both opportunities and challenges. While the framework's mathematical precision offers advantages for legal implementation, it also requires significant adaptations to existing legal doctrines and procedures.

Proportional Liability Systems

The GRN framework's continuous responsibility coefficients align well with legal systems that support proportional liability, such as comparative negligence doctrines. Courts could use GRN calculations to apportion damages among multiple defendants based on their respective responsibility coefficients. This approach would provide more precise and fair damage

allocation than current approaches, which often rely on rough judicial estimates of relative fault.

For implementation, courts would need access to expert testimony on GRN calculations and standardized methodologies for computing responsibility coefficients. This might require the development of new forms of legal expertise, similar to how courts have adapted to handle complex financial or scientific evidence in other domains.

Temporal Liability Adjustments

The temporal decay functions in GRN suggest that legal liability should be adjusted based on the time elapsed since different actors' involvement in AI systems. This could be implemented through modified statutes of limitations that account for the distributed and evolving nature of AI system responsibility.

Rather than uniform limitation periods, legal systems could adopt "responsibility-weighted" limitation periods where the time limits for bringing claims against different actors vary based on their responsibility coefficients and temporal decay patterns. Actors with higher initial responsibility coefficients might face longer limitation periods, while those whose responsibility decays quickly might benefit from shorter periods.

Emergent Behavior Liability

The GRN framework's approach to emergent behaviors suggests new legal doctrines for handling unintended AI system behaviors. Traditional legal frameworks struggle with emergent behaviors because they rely heavily on foreseeability standards. The GRN approach suggests a more nuanced framework that considers factors like the predictability of emergence, the magnitude of deviation from intended behavior, and the adequacy of response to emergent behaviors.

This could be implemented through "emergent behavior liability" doctrines that impose different standards of care for different types of actors based on their roles in creating conditions for emergence and their capacity to respond to emergent behaviors when they occur.

6.3 Corporate Governance and Risk Management

The GRN framework has significant implications for how organizations approach AI governance and risk management. The framework's emphasis on distributed responsibility suggests that effective AI governance requires coordination across multiple organizational functions and external stakeholders.

Responsibility-Based Governance Structures

Organizations developing or deploying AI systems could use GRN principles to design governance structures that align authority and accountability with responsibility coefficients. This might involve creating "responsibility matrices" that map different organizational roles to their expected responsibility coefficients for different types of AI outcomes.

For example, a technology company might assign primary responsibility for AI safety to its chief technology officer (high responsibility coefficient for technical decisions) while assigning primary responsibility for deployment decisions to business unit leaders (high responsibility coefficient for market and use case decisions). The governance structure would then ensure that decision-making authority and accountability mechanisms align with these responsibility distributions.

Dynamic Risk Assessment

The temporal aspects of GRN suggest that organizational risk assessments should be dynamic rather than static. Organizations should regularly recalculate responsibility coefficients as AI systems evolve, contexts change, and new stakeholders become involved. This dynamic approach would enable more responsive risk management and more effective allocation of risk mitigation resources.

Organizations could implement "responsibility monitoring systems" that track changes in responsibility coefficients over time and trigger governance interventions when coefficients exceed predetermined thresholds. For instance, if a deploying organization's responsibility coefficient for a particular AI system increases due to accumulating evidence of bias or safety issues, this could trigger enhanced oversight requirements or system modifications.

Network-Based Due Diligence

The network structure of GRN suggests that organizations should conduct due diligence not just on their direct AI vendors and partners, but on the broader networks of actors involved in AI systems. This "network due diligence" would involve understanding the responsibility distributions across entire AI ecosystems and identifying potential points of failure or inadequate accountability.

For example, a hospital deploying a medical AI system would conduct due diligence not only on the AI vendor, but also on the data providers, infrastructure providers, and other actors whose actions could affect the system's safety and effectiveness. This broader perspective would enable more comprehensive risk management and more effective accountability mechanisms.

6.4 Technical Standards and Best Practices

The GRN framework suggests several directions for developing technical standards and best practices for AI development and deployment. The framework's mathematical precision enables the development of quantitative standards that can guide technical decision-making and enable systematic comparison across different approaches.

Responsibility-Aware System Design

AI system designers could use GRN principles to create systems that facilitate appropriate responsibility attribution. This might involve designing systems with better logging and auditability features that enable post-hoc analysis of responsibility distributions. It could also involve designing systems with explicit responsibility interfaces that make clear how different actors' decisions contribute to system outcomes.

For example, AI systems could include "responsibility dashboards" that provide real-time information about how different actors' decisions are affecting system behavior and outcomes. These dashboards could help users understand their own responsibility for system outcomes and make more informed decisions about system use.

Standardized Responsibility Metrics

The GRN framework enables the development of standardized metrics for measuring and comparing responsibility across different AI systems and contexts. Industry organizations could develop standard methodologies for calculating responsibility coefficients, similar to how financial industries have developed standardized risk metrics.

These standardized metrics could facilitate better comparison and benchmarking across different AI systems and organizations. They could also enable the development of "responsibility ratings" for AI systems, similar to credit ratings or safety ratings, that would help users and deployers make more informed decisions.

Temporal Responsibility Tracking

The temporal aspects of GRN suggest the need for technical systems that can track how responsibility evolves over time. This might involve developing "responsibility ledgers" that maintain records of how different actors' contributions to AI systems change over time. These ledgers could use blockchain or other distributed ledger technologies to ensure transparency and immutability.

Such systems would enable more accurate post-hoc analysis of responsibility for AI outcomes and could support legal and regulatory processes that require understanding of temporal responsibility patterns.

6.5 Educational and Professional Development Implications

The adoption of GRN principles would require significant changes in how professionals are educated and trained for AI-related roles. The framework's emphasis on distributed responsibility suggests that all actors in AI ecosystems need better understanding of their roles and obligations.

Interdisciplinary AI Ethics Education

Current AI education programs typically focus on technical skills with limited attention to ethical and legal considerations. The GRN framework suggests the need for more interdisciplinary approaches that help technical professionals understand their roles in broader responsibility networks.

This might involve developing new curricula that combine technical AI training with ethics, law, and policy education. Professionals would learn not only how to build AI systems, but how to understand and manage their responsibility for system outcomes across different contexts and time periods.

Responsibility-Aware Professional Standards

Professional organizations in AI-related fields could develop standards and codes of conduct based on GRN principles. These standards would help professionals understand their responsibilities and provide guidance for ethical decision-making in complex AI contexts.

For example, data science professional organizations could develop standards that specify how data scientists should consider their responsibility coefficients when making decisions about data collection, model training, and system deployment. These standards could include specific guidance on how to assess and mitigate responsibility risks.

Continuing Education and Adaptation

The dynamic nature of AI systems and responsibility suggests the need for ongoing professional development that helps practitioners adapt to changing responsibility landscapes. Professional organizations could develop continuing education programs that help practitioners understand how their responsibilities evolve as AI systems and contexts change.

6.6 International Cooperation and Governance

The global nature of AI development and deployment creates challenges for implementing responsibility frameworks that require coordination across different legal and regulatory

systems. The GRN framework's mathematical precision and systematic approach could facilitate international cooperation on AI governance.

Harmonized Responsibility Standards

International organizations could use GRN principles to develop harmonized standards for AI responsibility that can be adapted to different legal and cultural contexts. The framework's mathematical foundation provides a common language for discussing responsibility that transcends specific legal traditions.

This could involve developing international agreements on responsibility calculation methodologies, similar to international accounting standards or environmental measurement protocols. Such agreements would facilitate cross-border AI governance and reduce regulatory arbitrage.

Cross-Border Responsibility Tracking

The network structure of GRN is well-suited to handling cross-border AI systems where different actors operate under different legal jurisdictions. The framework can track responsibility across jurisdictional boundaries and provide tools for coordinating governance responses.

International organizations could develop systems for sharing responsibility information across borders, enabling more effective coordination of regulatory responses to AI incidents that involve multiple jurisdictions.

6.7 Societal and Democratic Implications

The adoption of GRN principles has broader implications for how societies approach questions of technological governance and democratic accountability. The framework's emphasis on distributed responsibility aligns with democratic values of shared governance and collective decision-making.

Public Participation in AI Governance

The GRN framework's recognition that multiple stakeholders bear responsibility for AI outcomes suggests that governance processes should involve broader public participation. Citizens and civil society organizations are part of the networks affected by AI systems and should have roles in governance processes proportionate to their stakes in outcomes.

This could involve developing new forms of participatory governance that enable public input on responsibility allocation and accountability mechanisms. Citizens could participate in processes for setting responsibility standards and evaluating the performance of accountability systems.

Transparency and Democratic Oversight

The mathematical precision of GRN enables new forms of transparency and democratic oversight of AI systems. Responsibility coefficients and network analyses could be made public, enabling citizens and civil society organizations to understand and evaluate how responsibility is distributed across AI ecosystems.

This transparency could support more informed democratic debate about AI governance and enable more effective advocacy for accountability reforms. Citizens would have better tools for understanding how AI systems affect them and for holding relevant actors accountable for system outcomes.

The implications and applications of the GRN framework extend far beyond technical considerations to encompass fundamental questions about how societies organize themselves around powerful technologies. The framework provides tools for more democratic, transparent, and effective approaches to AI governance that can help ensure that the benefits of AI are broadly shared while minimizing potential harms. As AI systems become increasingly central to social and economic life, the need for such sophisticated approaches to accountability becomes ever more urgent.

7. Limitations and Future Research

7.1 Theoretical Limitations

While the Gradient Responsibility Networks framework represents a significant advance over existing approaches to AI moral accountability, it is important to acknowledge several theoretical limitations that constrain its applicability and suggest directions for future research.

Measurement Challenges

The mathematical precision of GRN depends on the ability to accurately measure the various factors that contribute to responsibility coefficients. However, some of these factors—particularly foreseeability, intentionality, and mitigation effort—involve subjective judgments that may be difficult to quantify consistently across different contexts and evaluators. The framework provides guidance for these assessments, but significant inter-rater reliability challenges may remain.

Future research should focus on developing more objective and standardized methodologies for measuring responsibility factors. This might involve creating detailed rubrics for assessing different types of foreseeability, developing behavioral indicators for intentionality, or creating standardized metrics for evaluating mitigation efforts. Machine learning approaches might also be developed to assist in these assessments by identifying patterns in how different factors correlate with responsibility attributions.

Cultural and Contextual Variations

The GRN framework is grounded in Western philosophical traditions of individual moral responsibility and may not translate directly to cultural contexts with different conceptions of agency, responsibility, and accountability. Some cultures emphasize collective responsibility or hierarchical accountability structures that may not align well with the framework's network-based approach.

Future research should explore how GRN principles can be adapted to different cultural contexts while maintaining their core analytical power. This might involve developing culturally-specific parameter settings for responsibility calculations or creating alternative formulations that better reflect different cultural values and social structures.

Complexity and Scalability

While the GRN framework is designed to handle complex networks of actors, there may be practical limits to its scalability. Very large AI systems involving thousands of contributors across multiple organizations and jurisdictions may create networks too complex for practical analysis. The computational requirements for calculating responsibility coefficients across very large networks may also become prohibitive.

Future research should explore methods for simplifying GRN analysis in very large systems while maintaining accuracy. This might involve developing hierarchical approaches that analyze responsibility at different levels of granularity or creating sampling methods that can estimate network-level responsibility patterns without analyzing every individual actor.

7.2 Empirical Limitations

The empirical validation of the GRN framework presented in this thesis, while comprehensive, has several limitations that suggest the need for additional research and validation.

Limited Case Study Diversity

The case studies presented focus primarily on high-profile AI incidents in developed Western contexts. Additional validation is needed across a broader range of AI applications, cultural contexts, and types of harm. The framework's performance may vary significantly across different domains such as AI in education, environmental monitoring, or social services.

Future research should systematically apply the GRN framework across a broader range of case studies, including cases from different cultural contexts, different types of AI systems, and different scales of impact. Longitudinal studies tracking how responsibility attributions

change over time would be particularly valuable for validating the framework's temporal components.

Lack of Experimental Validation

The current validation relies primarily on post-hoc analysis of existing cases rather than experimental testing of the framework's predictive power. It remains unclear how well GRN-based responsibility attributions align with human moral intuitions or how effectively the framework can guide prospective decision-making.

Future research should include experimental studies that test how GRN-based responsibility attributions compare with human moral judgments across different scenarios. These studies could also test whether GRN-based guidance improves decision-making outcomes compared to alternative approaches.

Limited Stakeholder Input

The development and validation of the GRN framework has involved limited input from key stakeholders such as AI practitioners, legal professionals, policymakers, and affected communities. The framework's practical utility and acceptability may depend significantly on stakeholder buy-in and adaptation to their needs and constraints.

Future research should involve extensive stakeholder engagement to refine the framework based on practical experience and to develop implementation strategies that address stakeholder concerns and constraints.

7.3 Implementation Limitations

The practical implementation of GRN principles faces several challenges that limit its immediate applicability and suggest the need for significant institutional adaptation.

Legal System Integration

While the GRN framework is designed to be compatible with existing legal systems, its integration would require significant changes to legal procedures, evidence standards, and judicial training. Courts would need to develop expertise in evaluating GRN-based analyses and determining appropriate weight to give such evidence.

The framework's mathematical complexity may also create barriers to access for parties who cannot afford expert analysis. This could exacerbate existing inequalities in legal system access and create new forms of disadvantage for less resourced parties.

Regulatory Capacity

The implementation of GRN-based regulatory approaches would require significant investments in regulatory capacity and expertise. Regulatory agencies would need to develop new forms of technical expertise and new institutional mechanisms for network-based oversight.

The dynamic nature of GRN-based regulation would also require more flexible and adaptive regulatory approaches than many agencies currently employ. This institutional adaptation may face significant political and bureaucratic obstacles.

Industry Resistance

The implementation of GRN principles may face resistance from AI industry actors who prefer simpler and more predictable accountability frameworks. The framework's emphasis on distributed responsibility may create concerns about expanded liability exposure and increased compliance costs.

Successful implementation will likely require careful attention to industry concerns and the development of implementation strategies that balance accountability goals with innovation incentives.

7.4 Future Research Directions

The limitations identified above suggest several important directions for future research that could strengthen the theoretical foundations, empirical validation, and practical implementation of the GRN framework.

Advanced Mathematical Modeling

Future research should explore more sophisticated mathematical approaches to responsibility modeling. This might include developing dynamic network models that can capture how responsibility networks evolve over time, exploring machine learning approaches to responsibility coefficient estimation, or developing game-theoretic models of how different actors' responsibility calculations interact.

Research on uncertainty quantification would also be valuable, providing tools for understanding and communicating the confidence intervals around responsibility coefficient estimates. This would help decision-makers understand the limitations of GRN-based analyses and make more informed judgments about their reliability.

Behavioral and Psychological Research

Understanding how the GRN framework aligns with human moral psychology is crucial for its acceptance and effectiveness. Future research should explore how people naturally think about responsibility in complex AI contexts and how GRN-based framings affect moral judgments and decision-making.

Research on the behavioral effects of different responsibility attribution approaches would also be valuable. Do GRN-based approaches to accountability create better incentives for responsible AI development and deployment than alternative approaches? How do different actors respond to GRN-based responsibility assessments?

Comparative Legal Analysis

Future research should explore how GRN principles can be integrated into different legal systems and traditions. Comparative analysis of how different legal systems handle complex responsibility attribution could identify best practices and potential obstacles to GRN implementation.

Research on the development of new legal doctrines and procedures specifically designed for GRN-based analysis would also be valuable. This might include developing new forms of expert testimony, new approaches to damage calculation, or new procedural mechanisms for handling complex responsibility networks.

Technology-Assisted Implementation

The complexity of GRN analysis suggests the need for technological tools that can assist in responsibility calculation and network analysis. Future research should explore the development of software tools that can automate aspects of GRN analysis while maintaining transparency and accountability.

Research on blockchain and other distributed ledger technologies for responsibility tracking could also be valuable, providing tools for maintaining transparent and immutable records of how responsibility evolves over time.

7.5 Broader Implications for AI Governance

The development and refinement of the GRN framework is part of a broader challenge of adapting governance institutions to the realities of AI and other emerging technologies. The limitations identified above reflect not just shortcomings of the specific framework, but broader challenges in governing complex technological systems.

Institutional Innovation

The implementation of sophisticated approaches to AI accountability like GRN will likely require significant institutional innovation. Traditional governance institutions—courts,

regulatory agencies, professional organizations—may need fundamental adaptations to handle the complexity and dynamism of AI systems.

Future research should explore what new institutional forms might be needed for effective AI governance and how existing institutions can be adapted to handle new challenges. This might include research on new forms of multi-stakeholder governance, new approaches to international coordination, or new mechanisms for public participation in technical decision-making.

Democratic Governance of Technology

The GRN framework raises broader questions about how democratic societies can maintain meaningful control over powerful technologies. The framework's emphasis on distributed responsibility aligns with democratic values, but its technical complexity may create barriers to public understanding and participation.

Future research should explore how technical frameworks like GRN can be made more accessible to public understanding and how they can support rather than undermine democratic governance of technology. This might include research on new forms of science communication, new approaches to public participation in technical decision-making, or new mechanisms for democratic oversight of algorithmic systems.

8. Conclusion

This thesis has examined one of the most pressing questions in contemporary technology policy: Should we be morally accountable for AI behavior? Through comprehensive analysis of philosophical foundations, empirical case studies, and practical applications, the research provides a definitive answer: Yes, we should be morally accountable for AI behavior, but this accountability must be understood as distributed across networks of actors rather than concentrated in individual agents.

8.1 Key Findings and Contributions

The research makes several significant contributions to the growing field of AI ethics and governance:

Theoretical Innovation: The Gradient Responsibility Networks framework represents a fundamental advance in thinking about moral accountability in complex technological systems. By treating responsibility as continuous rather than binary, incorporating temporal dynamics, and providing tools for analyzing network-based responsibility distributions, GRN addresses critical gaps in existing approaches that have limited their effectiveness in AI contexts.

Empirical Validation: The application of GRN principles to real-world case studies—including the Tesla Autopilot crashes, Amazon's hiring algorithm bias, and medical AI diagnostic errors—demonstrates the framework's practical applicability and superior performance compared to traditional approaches. The case studies reveal consistent patterns of distributed responsibility that validate the framework's theoretical foundations.

Practical Implementation: Unlike purely theoretical frameworks, GRN is designed for practical implementation in legal, regulatory, and organizational contexts. The framework's mathematical precision enables quantitative analysis while its network-based structure provides tools for understanding complex stakeholder relationships.

Comparative Advantage: Systematic comparison across six critical dimensions demonstrates GRN's substantial performance improvements over existing approaches, with improvements ranging from 75% to 350% across different evaluation criteria.

8.2 Implications for AI Governance

The research has profound implications for how societies approach AI governance and accountability:

Regulatory Reform: The GRN framework suggests the need for dynamic, network-based regulatory approaches that can adapt to the evolving nature of AI systems. Traditional static regulatory frameworks are inadequate for the temporal and distributed characteristics of AI accountability.

Legal System Adaptation: The integration of GRN principles into legal frameworks would require significant adaptations to existing doctrines and procedures, but would provide more precise and fair approaches to liability attribution in complex AI cases.

Organizational Governance: The framework provides tools for organizations to design governance structures that align authority and accountability with actual responsibility distributions, enabling more effective risk management and ethical decision-making.

Democratic Participation: The framework's emphasis on distributed responsibility suggests the need for broader public participation in AI governance, with citizens and civil society organizations playing roles proportionate to their stakes in AI outcomes.

8.3 Addressing the Central Question

Returning to the central question posed in the thesis title—Should we be morally accountable for AI behavior?—the research provides a nuanced but definitive answer. We should indeed be morally accountable for AI behavior, but this accountability must be understood through frameworks sophisticated enough to capture the complex realities of modern AI systems.

The traditional approach of seeking single responsible parties for AI outcomes is not merely inadequate—it is counterproductive. By forcing artificial binary choices about responsibility, traditional frameworks often result in either over-attribution (holding one party responsible for outcomes involving multiple contributors) or under-attribution (failing to hold anyone responsible when no single party bears complete responsibility). Both outcomes undermine the goals of accountability: preventing harm, ensuring justice, and promoting beneficial technological development.

The GRN framework demonstrates that distributed accountability is not a dilution of responsibility but rather a more accurate reflection of how AI systems actually work. When multiple actors contribute to AI outcomes through complex interactions over extended time periods, accountability frameworks must be sophisticated enough to capture these relationships. The framework's ability to generate total network responsibility values greater than 1.0 reflects the reality that collective responsibility can exceed the sum of individual contributions—a phenomenon that traditional frameworks cannot capture.

8.4 The Urgency of Reform

The empirical evidence presented in this thesis underscores the urgency of developing better approaches to AI accountability. With AI project failure rates exceeding 80% and documented bias incidents across multiple sectors causing significant harm to individuals and communities, the status quo is clearly inadequate. The rapid advancement of AI capabilities and their deployment in increasingly critical domains makes the development of effective accountability frameworks not merely important but essential for the continued beneficial development of AI technologies.

The case studies examined—from fatal autonomous vehicle crashes to discriminatory hiring algorithms to biased medical diagnostic tools—represent only the visible tip of a much larger iceberg of AI-related harms. As AI systems become more powerful and ubiquitous, the potential for both beneficial and harmful outcomes increases exponentially. Without appropriate accountability frameworks, societies risk losing control over these powerful technologies and failing to ensure that their benefits are broadly shared while their harms are minimized.

8.5 A Call for Action

The research presented in this thesis provides both a diagnosis of current problems and a prescription for solutions. The GRN framework offers concrete tools for improving AI accountability, but its implementation will require coordinated action across multiple domains:

Academic Research: Continued research is needed to refine the framework, validate its performance across broader contexts, and develop supporting tools and methodologies. The interdisciplinary nature of AI accountability requires collaboration across computer science, philosophy, law, policy, and social sciences.

Policy Development: Policymakers and regulators need to begin adapting existing frameworks and developing new approaches based on GRN principles. This will require significant investment in regulatory capacity and expertise, as well as willingness to experiment with new approaches to governance.

Industry Adoption: Technology companies and other organizations developing and deploying AI systems need to begin implementing GRN-based approaches to governance and risk management. This will require changes to organizational structures, decision-making processes, and accountability mechanisms.

Legal System Evolution: Legal systems need to begin adapting to handle the complexity of AI accountability through new doctrines, procedures, and forms of expertise. This will require training for legal professionals and development of new institutional mechanisms.

Public Engagement: Citizens and civil society organizations need to become more engaged in AI governance processes, both to ensure democratic accountability and to provide essential perspectives on the impacts of AI systems on different communities.

8.6 Looking Forward

The question of moral accountability for AI behavior will only become more pressing as AI systems become more powerful and autonomous. The frameworks we develop today will shape the trajectory of AI development for decades to come. The choice is not whether to hold actors

accountable for AI behavior—the choice is whether to do so through sophisticated frameworks that capture the complex realities of AI systems or through inadequate frameworks that fail to serve the goals of accountability.

The Gradient Responsibility Networks framework provides a foundation for more effective and just approaches to AI accountability. However, it is not a final answer but rather a starting point for ongoing research, development, and implementation. The framework will need to be refined based on practical experience, adapted to different contexts and cultures, and integrated with other approaches to AI governance.

The ultimate goal is not perfect accountability—an impossible standard in any complex domain—but rather accountability systems that are sophisticated enough to handle the challenges posed by AI while remaining practical enough for real-world implementation. The GRN framework provides tools for moving toward this goal, but achieving it will require sustained effort and commitment from all stakeholders in the AI ecosystem.

8.7 Final Reflections

The development of artificial intelligence represents one of the most significant technological advances in human history, with the potential to transform virtually every aspect of human life. With this tremendous potential comes tremendous responsibility—not just for the developers and deployers of AI systems, but for all of us who shape the contexts in which these systems operate and who are affected by their outcomes.

The question of moral accountability for AI behavior is ultimately a question about what kind of society we want to live in and what kind of future we want to create. Do we want a future where powerful AI systems operate without adequate accountability, where harms go unaddressed and benefits are unevenly distributed? Or do we want a future where sophisticated accountability frameworks ensure that AI systems serve human flourishing and reflect our deepest values?

The research presented in this thesis argues strongly for the latter vision. By developing and implementing frameworks like GRN that can capture the distributed, temporal, and emergent nature of AI systems, we can work toward a future where AI serves humanity rather than the

reverse. This is not just a technical challenge but a moral imperative—one that requires the best of our intellectual, institutional, and ethical capabilities.

The stakes could not be higher, and the time for action is now. The frameworks we develop today for AI accountability will determine whether artificial intelligence becomes a force for human flourishing or a source of unprecedented harm. The choice is ours, and the responsibility—distributed though it may be—is real.

References

- [1] National Highway Traffic Safety Administration. (2017). Investigation PE 16-007: Tesla Model S and Model X Automatic Vehicle Control Systems. NHTSA Office of Defects Investigation. <https://static.nhtsa.gov/odi/inv/2016/INCLA-PE16007-7876.PDF>
- [2] RAND Corporation. (2024). The Root Causes of Failure for Artificial Intelligence Projects and How They Can Succeed. RAND Research Reports, RRA2680-1. https://www.rand.org/pubs/research_reports/RRA2680-1.html
- [3] S&P Global Market Intelligence. (2025). AI Project Failure Rates Continue to Rise. Technology Research Report. <https://www.spglobal.com/marketintelligence/en/news-insights/research/ai-project-failures-2025>
- [4] Aristotle. (4th century BCE). Nicomachean Ethics. Translated by W.D. Ross. Oxford: Oxford University Press, 1998.
- [5] McKenna, M. & Russell, P. (2024). Moral Responsibility. Stanford Encyclopedia of Philosophy. <https://plato.stanford.edu/entries/moral-responsibility/>
- [6] van Inwagen, P. (1983). An Essay on Free Will. Oxford: Oxford University Press.
- [7] Gilbert, M. (2000). Sociality and Responsibility: New Essays in Plural Subject Theory. Lanham, MD: Rowman & Littlefield.
- [8] Stanford Encyclopedia of Philosophy. (2024). Moral Responsibility. <https://plato.stanford.edu/entries/moral-responsibility/>
- [9] Martinho, A., Herber, N., Kroesen, M., & Chorus, C. (2021). Ethical issues in focus: Perspectives on artificial moral agents. *AI & Society*, 36(3), 719-734. <https://link.springer.com/article/10.1007/s00146-020-01020-6>

- [10] Wallach, W., & Allen, C. (2008). *Moral Machines: Teaching Robots Right from Wrong*. Oxford: Oxford University Press.
- [11] Novelli, C., Taddeo, M., & Floridi, L. (2023). Accountability in artificial intelligence: what it is and how it works. *AI & Society*, 39, 1871–1882.
<https://link.springer.com/article/10.1007/s00146-023-01635-y>
- [12] European Union. (2024). Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act). Official Journal of the European Union.
<https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:32024R1689>
- [13] RAND Corporation. (2024). Liability for Harms from AI Systems. Research Report RRA3243-4. https://www.rand.org/pubs/research_reports/RRA3243-4.html
- [14] O'Neil, C. (2016). *Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy*. New York: Crown Publishers.
- [15] MIT Technology Review. (2024). Why 85% of AI Projects Fail.
<https://www.technologyreview.com/2024/07/30/ai-project-failure-rates/>
- [16] S&P Global Market Intelligence. (2025). AI Initiative Abandonment Rates Surge. Market Intelligence Report.
<https://www.spglobal.com/marketintelligence/en/news-insights/research/ai-abandonment-2025>
- [17] NTT Data. (2024). GenAI Deployment Failure Analysis. Technology Research Report.
<https://www.nttdata.com/global/en/insights/focus/2024/genai-deployment-failures>
- [18] Dynatrace. (2024). Why 85% of AI Projects Fail and How to Save Yours. Technical Report. <https://www.dynatrace.com/news/blog/why-ai-projects-fail/>
- [19] Wall Street Journal. (2024). The Hidden Autopilot Data That Reveals Why Teslas Crash. Investigative Report. <https://www.wsj.com/articles/tesla-autopilot-crash-data-analysis>

[20] Dastin, J. (2018). Amazon scraps secret AI recruiting tool that showed bias against women. Reuters. <https://www.reuters.com/article/world/insight-amazon-scraps-secret-ai-recruiting-tool-that-showed-bias-against-women-idUSKCN1MK0AG/>

[21] European Commission. (2024). The AI Act: Europe's Answer to AI Governance. Policy Brief. <https://digital-strategy.ec.europa.eu/en/policies/regulatory-framework-ai>